IAC–17–C1,4,2,x41662

Comparison of EKF and UKF for Robotic Missions to Mars

Jeffrey Stuart
Jet Propulsion Laboratory, California Institute of Technology, USA, jeffrey.r.stuart@jpl.nasa.gov

Todd Ely
Jet Propulsion Laboratory, California Institute of Technology, USA, todd.a.ely@jpl.nasa.gov

The batched Kalman filter works well for a wide variety of orbit determination scenarios but requires linearized approximations to the underlying dynamical system. Furthermore, implementation depends on the computation of partial derivatives for all measurements and parameters used in the filter, a non-trivial task even when analytical derivations are possible. To address these limitations, many extensions and modifications to the linear Kalman filter have been proposed over the years since Kalman's initial publication. Our investigation examines two popular variants of the Kalman filter, namely the Extended Kalman Filter (EKF) and Unscented Kalman Filter (UKF), within the context of spacecraft missions to Mars. The EKF relies on the same mathematical basis as the classic version of the filter, however it sequentially updates the linearization as the filtering process is conducted. In contrast, the UKF is built upon the unscented transform which samples the nearby solution space and numerically propagates perturbed trajectories as part of the update / estimation process. We assess the performance of the filtering algorithms under a variety of mission scenarios with the goal of establishing decision criteria for which implementation to use under differing circumstances.

I. Introduction

Orbit determination, or the process of turning indirect measurements of a spacecraft's motion into an updated understanding its state and the associated uncertainties, is one of the key operational requirements for deep space exploration. Since its initial publication in 1960, the linear Kalman filter[1] has grown to be the main orbit determination algorithm for companies and space agencies across the world and has proven its versatility in supporting planetary exploration across the solar system. In particular, the batched variant wherein all measurements are processed simultaneously rather than sequentially has become an embedded part of most operational spaceflight missions. Indeed, it is hard to overstate the importance and utility of the Kalman filter to many applications across the aerospace industry.[2–4] However, the traditional Kalman filter has two key requirements: (i) the system considered must be sufficiently close to linear; and, (ii) partial derivatives of the states and measurements must be derived and integrated along with the equations of motion. To address the first deficiency, methods of linearizing orbital motion relative to a reference path are usually employed. Since these linearizations degrade in accuracy over time, navigators have implemented work flows wherein spacecraft state estimates are updated on a rolling basis as new tracking passes are conducted. A popular modification of the traditional linear Kalman filter (LKF) to nonlinear systems is the extended Kalman filter (EKF), though most applications make use of the sequential version wherein measurements are processed into updated estimates as the measurement arrive.[5] While often used for spaceflight applications, the sequential EKF presents additional challenges when trying to estimate short-duration events such as maneuvers, at least when continuous tracking is not assumed. Furthermore, the EKF still relies on the provision of state and measurement partials to the estimation program.

Over the past few decades, alternatives to the Kalman filter have been gaining in popularity, both in academic research and in industrial applications; two exemplars of this trend are the particle filter[6] and the unscented Kalman filter (UKF).[7–9] These more modern filters have the notable benefit that they do not require the use of partial derivatives or the state transition matrix, however this comes at the cost of an increased computational overhead because virtual particles must be propagated alongside the main motion of interest. This trade-off limited the use of these classes of filters until the common availability of parallel computing capabilities such as multiple cores, dedicated clusters, GPUs, and cloud computing; now, however, the UKF and other filters are gaining in

popularity across a variety of industries. Within the field of spaceflight, the UKF has been variously examined for use in: space-object tracking;[10] covariance assessment[11] and real-time orbit determination[12] using GPS signals; formation flying relative positioning;[13] optical navigation for rendezvous[14] and small body exploration;[15] and lunar as well as planetary exploration.[16] However, all of these investigations focused on use of the sequential UKF, a useful implementation for real-time applications such as on-board autonomy but not strictly required for ground-in-the-loop operations.

In our investigation, we apply batch versions of the LKF,[4] EKF,[17] and UKF[18] to robotic spaceflight missions to Mars. Our goal is to determine what performance gains, if any, can come from use the batched EKF or UKF and what operational scenarios might be enabled by these newer filter methods. We consider different phases of a potential Mars mission, from approach to the Mars system from interplanetary space down to operations in a low Mars orbit. For the moment, we assume human-in-the-loop operation, so we assess the filter performance terms of resources that would be available on the ground. In addition to updating the spacecraft state, we have also modified the batch UKF to enable estimation of the state uncertainty over time, with a particular emphasis on forward predictions of the covariance ellipsoids. To simplify our analysis, we have included only use of the two-way Doppler radiometric measurements, though other measurement types could be readily included. We performed all trajectory and navigation simulations using JPL's Mission Analysis, Operations, and Navigation Toolkit Environment (MONTE);[19] planetary ephemerides are provided by JPL's HORIZONS database.[20]

II. Mathematical Formulations

We begin by describing the mathematical basis for the various Kalman filtering algorithms and their implementation as batch-filtering methods for orbit determination. We omit mathematical derivations of orbital mechanics and the Doppler measurement type since many other sources provide detailed examinations of these topics.[21–23] The goal of orbit determination is to update estimates of a set of M parameters $\boldsymbol{x}$ subject to the dynamical equation

$$\boldsymbol{x}_{k+1} = \boldsymbol{f}(\boldsymbol{x}_k, \mathbf{w}_k, t_k) \quad [1]$$

where $\boldsymbol{f}$ describes the system dynamics, $\boldsymbol{x}_k$ represents the parameter values at reference epoch t_k, and $\mathbf{w}_k$ is the unmodeled process noise. Note that the parameter vector $\boldsymbol{x}$ can contain both dynamic (time-varying) and bias (time-invariant) parameters. For example, position and velocity are always included as dynamic parameters, while maneuvers, spacecraft shapes, and gravitational models are often included as bias parameters. By definition, the time derivatives of bias parameters are zero; for bias terms that may change over time but for which we can't adequately model all components of the dynamics (e.g., solar radiation pressure, atmospheric drag), the standard practice is to include stochastic bias terms over shorter segments of the trajectory. The parameter vector $\boldsymbol{x}_k$ has a corresponding covariance $\boldsymbol{P}_k$, where estimates of both are updated according to the measurement set

$$\boldsymbol{y}_k = \boldsymbol{h}(\boldsymbol{x}_k, t_k) + \mathbf{v}_k \quad [2]$$

where the vector $\boldsymbol{y}_k$ is a collection of N measurements taken over the interval $t_1 \leq t \leq t_N$, $\boldsymbol{h}$ represents the mapping of parameters to measurements either through computation or observation and $\mathbf{v}_k$ is the measurement noise. The noises $\mathbf{w}_k$ and $\mathbf{v}_k$ are usually assumed to be uncorrelated, zero-mean Gaussian distributions with the respective covariances $E\{\mathbf{w}_k\mathbf{w}_k^T\} = \boldsymbol{Q}_k$ and $E\{\mathbf{v}_k\mathbf{v}_k^T\} = \boldsymbol{R}_k$. Note that the general formulation, as currently posed, permits nonlinear and even discontinuous dynamic terms $\boldsymbol{f}$ and measurements $\boldsymbol{h}$, however practical filter implementation often poses stricter requirements.

II.i Linear Kalman Filter (LKF)

In this section, we present the mathematical operations composing the batch linear Kalman filter (LKF), also known as the linear least squares estimator; we follow the derivation in Ch. 4.4 of *Statistical Orbit Determination*.[4] The batch LKF assumes linear mapping in the parameters $\boldsymbol{x}_k$ from the reference epoch t_k. That is, the LKF assumes that the dynamics represented by Eq. [1] can be linearized to

$$\delta\dot{\boldsymbol{x}}_k = \boldsymbol{F}_k\delta\boldsymbol{x}_k + \mathbf{w}_k \quad [3]$$

and the measurements in Eq. [2] are likewise linearized as

$$\delta\boldsymbol{y}_k = \boldsymbol{H}_k\delta\boldsymbol{x}_k + \mathbf{v}_k \quad [4]$$

where $\boldsymbol{F}_k$ and $\boldsymbol{H}_k$ are the partial derivatives of the dynamic and measurement mappings with respect to the parameters $\boldsymbol{x}$. The term $\delta\boldsymbol{x}_k$ denotes a variation from the reference set of parameters. For problems in orbit determination, satisfying the assumption of linearity necessitates at least one of three criteria be met:

- the variation $\delta \boldsymbol{x}_k$ is sufficiently small;
- the time spans $[t_k, t_{k+1}]$ and $[t_1, t_N]$ are sufficiently small and overlapping; or,
- the dynamics are relatively uniform, e.g., the spacecraft is far from any gravitating bodies or other non-uniform perturbations.

As a practical matter, some combination of all three is usually imposed, depending on the exact circumstances; balancing these three considerations in spaceflight operations is as much art as science.

The batch LKF processes the complete set of N observed measurements $\tilde{\boldsymbol{y}}_k$ in order to update the parameter estimates $\hat{\boldsymbol{x}}_k$ through an iterative process. The parameter update is given by

$$\hat{\boldsymbol{x}}_k^+ = \boldsymbol{K}^{-1}\left(\boldsymbol{H}_k^T \boldsymbol{R}_k^{-1} \Delta\tilde{\boldsymbol{y}}_k + (\boldsymbol{P}_k^-)^{-1}\hat{\boldsymbol{x}}_k^-\right) \quad [5]$$

where the superscripts $(-,+)$ indicate pre- and post-update, respectively, and $\Delta\tilde{\boldsymbol{y}}_k = \tilde{\boldsymbol{y}}_k - \hat{\boldsymbol{y}}_k^-$ is the measurement residual (the difference between the observed measurements $\tilde{\boldsymbol{y}}_k$ and simulated measurements $\hat{\boldsymbol{y}}_k^-$). The Kalman gain $\boldsymbol{K}$ is computed via

$$\boldsymbol{K} = \boldsymbol{H}_k^T \boldsymbol{R}_k^{-1} \boldsymbol{H}_k + (\boldsymbol{P}_k^-)^{-1} \quad [6]$$

and the *a posteriori* covariance estimate is calculated as

$$\boldsymbol{P}_k^+ = \boldsymbol{K}^{-1}. \quad [7]$$

The measurement set is processed by summing over the individual measurements via

$$\boldsymbol{H}_k^T \boldsymbol{R}_k^{-1} \boldsymbol{H}_k = \sum_{i=1}^{N} (\boldsymbol{H}_i \boldsymbol{\Phi}_{i,k})^T \boldsymbol{R}_i^{-1} (\boldsymbol{H}_i \boldsymbol{\Phi}_{i,k}) \quad [8]$$

and

$$\boldsymbol{H}_k^T \boldsymbol{R}_k^{-1} \Delta\tilde{\boldsymbol{y}}_k = \sum_{i=1}^{N} (\boldsymbol{H}_i \boldsymbol{\Phi}_{i,k})^T \boldsymbol{R}_i^{-1} \Delta\tilde{\boldsymbol{y}}_i \quad [9]$$

where the subscript i indicates the time t_i as well as its associated measurement while $\boldsymbol{\Phi}_{i,k}$ is the state transition matrix mapping variations in $\boldsymbol{x}_k$ to corresponding changes in $\boldsymbol{x}_i$. In the batch LKF, all measurements are processed with respect to the estimated parameters at the singular reference epoch t_k, hindering the continued assumption of linearity and omitting higher order terms that may grow large over time. In addition, we are required to somehow derive and/or compute the partial derivatives of both the measurements and the parameters themselves. Finally, we note that the inversion of the Kalman gain is an unstable operation, leading to the common usage of UDU decompositions or square-root information formulations.[24]

Before proceeding to the other filtering methods, we wish to make a brief comment on stochastic bias parameters, even though we do not explicitly include them in this investigation. These variables are usually modeled as piecewise time series' applied over the course of the baseline trajectory: the total time span is divided into sub-intervals in which the stochastic parameter can take on different values within the different time bins. The LKF traditionally addresses these cases by adjusting the covariance of the stochastic parameters at the batch boundary using a variety of different noise models (random bias, white noise, random walk, etc.) which can either assume independence or correlation between the discrete sub-intervals. A smoothing process is then applied to map estimates across the full span of the trajectory (for more information, see Ch. 4.15 of *Statistical Orbit Determination*).[4]

II.ii Extended Kalman Filter (EKF)

The batch extended Kalman filter (EKF) is mathematically identical to the batch LKF, except that the full measurement set is partitioned into sub-batches which are processed in sequence, as illustrated in Fig. 1.[17] At each subset boundary, the parameter set

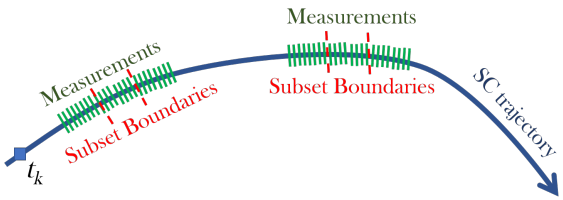

Fig. 1: Schema of EKF partitioning of full measurement set into smaller partitions.

is updated and reinitialized to $\hat{\boldsymbol{x}}_p$ and $\boldsymbol{P}_p$, where the subscript indicates the epoch t_p defining the partition between sub-batches. The motion is again linearized as in Eq. [3], but with respect to the intermediate epoch t_p, in preparation for processing the next subset of measurements. The process is analogous to the more common sequential EKF, though the sequential version processes measurements individually and updates the linearized reference motion at each measurement epoch. Once the complete batch of measurements is processed (i.e., all subsets are handled), the filtered solution is re-mapped to the epoch t_k and the entire process can be iterated as appropri-

ate. Stochastic parameters can be readily incorporated into the EKF using the same methods as the LKF, but as a practical manner the boundaries of the measurement and stochastic sub-batches should be aligned whenever possible. We note that, conceptually, the standard orbit determination workflow implicitly operates as a large-scale batch EKF wherein the long-term trajectory is split into segments and the spacecraft state and other parameters are estimated episodically.

II.iii Unscented Kalman Filter (UKF)

We now turn to the batch unscented Kalman filter (UKF) which exploits the unscented transform. Our formulation follows that of Park *et al.*,[18] though we extend the mathematical basis to include updated estimates of the parameter covariances if so desired and we use a different weighting scheme. We begin by initializing a symmetric set of sigma points $\bar{\boldsymbol{\chi}}_{j,k}$ for $j = 1, \ldots, M$ (recall that M is the number of parameters in $\boldsymbol{x}_k$):

$$\begin{aligned} \bar{\boldsymbol{\chi}}_{0,k} &= \hat{\boldsymbol{x}}_k \\ \bar{\boldsymbol{\chi}}_{j,k} &= \hat{\boldsymbol{x}}_k + \alpha\sqrt{\frac{M}{1-W^{(0)}}}(\boldsymbol{P}_k^{1/2})_{col\_j} \\ \bar{\boldsymbol{\chi}}_{j+M,k} &= \hat{\boldsymbol{x}}_k - \alpha\sqrt{\frac{M}{1-W^{(0)}}}(\boldsymbol{P}_k^{1/2})_{col\_j} \end{aligned} \quad (10)$$

where $\alpha \leq 1$ is a weighting parameter, $W^{(0)}$ is the weight of the base point (typically set to 0), and $(\boldsymbol{P}_k^{1/2})_{col\_j}$ is the j^{th} column of the Cholesky decomposition of the covariance matrix. Unless otherwise noted, this investigation uses a value of $\alpha = 5 \times 10^{-3}$ for all UKF runs. A corresponding set of sigma weights are calculated as

$$\begin{aligned} W_s^{(0)} &= \frac{W^{(0)} + \alpha^2 - 1}{\alpha^2} \\ W_s^{(j)} &= \frac{1 - W^{(0)}}{2M\alpha^2} \end{aligned} \quad (11)$$

where all weights sum to one. All $2M+1$ sigma points are then independently propagated across all measurement epochs and estimated measurements are computed for each sigma point as

$$\tilde{\boldsymbol{\gamma}}_{j,k} = \tilde{\boldsymbol{h}}(\bar{\boldsymbol{\chi}}_{j,k}) = \begin{bmatrix} h(\bar{\boldsymbol{\chi}}_{j,1}, t_1) \\ \vdots \\ h(\bar{\boldsymbol{\chi}}_{j,N}, t_N) \end{bmatrix} \quad [12]$$

for $j = 0, \ldots, 2M$ and $i = 1, \ldots, N$. The propagation of the sigma points and computation of the estimated measurement batches is readily parallelized for computational efficiency, though we have not done so in this investigation. However, note that the number of operations in the UKF ($2M + 1$ propagations of M parameters only) is of same order of magnitude as the LKF and EKF (propagation of both M parameters and M^2 partials), so the computational speed of a properly parallelized UKF approaches that of the standard LKF / EKF. Furthermore,

Once the sigma points and their associated measurement sets are calculated, they are combined to form estimates of the true measurements and the associated covariances. The predicted measurement is given by

$$\bar{\boldsymbol{y}}_k = \sum_{j=0}^{2M} W_s^{(j)} \tilde{\boldsymbol{\gamma}}_{j,k} \quad [13]$$

while the measurement covariance matrix is

$$\bar{\boldsymbol{P}}_k^{yy} = \sum_{j=0}^{2M} W_s^{(j)} (\tilde{\boldsymbol{\gamma}}_{j,k} - \bar{\boldsymbol{y}}_k)(\tilde{\boldsymbol{\gamma}}_{j,k} - \bar{\boldsymbol{y}}_k)^T + \boldsymbol{R}_k \quad [14]$$

where we recall that $\boldsymbol{R}_k$ is the uncorrelated measurement noise matrix. The parameter / measurement cross-correlation matrix is calculated

$$\bar{\boldsymbol{P}}_k^{xy} = \sum_{j=0}^{2M} W_s^{(j)} (\bar{\boldsymbol{\chi}}_{j,k} - \hat{\boldsymbol{x}}_k)(\tilde{\boldsymbol{\gamma}}_{j,k} - \bar{\boldsymbol{y}}_k)^T. \quad [15]$$

The Kalman gain for the batch UKF is given by

$$\boldsymbol{K}_k = \bar{\boldsymbol{P}}_k^{xy} (\bar{\boldsymbol{P}}_k^{yy})^{-1} \quad [16]$$

though as a computational matter the inversion is still unstable, so in practice LU-decompositions or least-squares equation solvers are used. The parameter estimate is updated by

$$\hat{\boldsymbol{x}}_k^+ = \hat{\boldsymbol{x}}_k^- + \boldsymbol{K}_k \Delta\tilde{\boldsymbol{y}}_k \quad [17]$$

where the measurement residual is $\Delta\tilde{\boldsymbol{y}}_k = \tilde{\boldsymbol{y}}_k - \bar{\boldsymbol{y}}_k$. As with the batch versions of the LKF and EKF, the whole process is iterated until convergence on a stable estimate or the measurement residuals assume a zero mean distribution. Note the advantage that the UKF does not rely on explicit derivation or computation of the parameter and measurement partials; because of this, discontinuities and non-smooth cases are easily handled without need for extensive modification. Furthermore, stochastic parameters are readily handled by including the full stochastic time series as a set of independent parameters with appropriate covariances defined for cross- and uncorrelated noise models. This does necessarily increase the size of the problem and the number of required sigma points,

however this can be mitigated to some extend by the parallelization.

The UKF can also update the parameter uncertainty or map it to another epoch t_{k+1} so long as the sigma point propagations extend to cover said epoch. The estimated parameters at the epoch t_{k+1} are calculated via

$$\bar{\boldsymbol{x}}_{k+1} = \sum_{j=0}^{2M} W_s^{(j)} \bar{\boldsymbol{\chi}}_{j,k+1} \quad [18]$$

with *a priori* parameter covariance

$$\bar{\boldsymbol{P}}_{k+1}^{-} = \sum_{j=0}^{2M} W_s^{(j)} (\bar{\boldsymbol{\chi}}_{j,k+1} - \hat{\boldsymbol{x}}_{k+1})(\bar{\boldsymbol{\chi}}_{j,k+1} - \hat{\boldsymbol{x}}_{k+1})^T + \boldsymbol{Q}_{k+1} \quad [19]$$

where $\boldsymbol{Q}_{k+1}$ is the small, uncorrelated process noise. The measurement data and covariance are computed as before, but now the cross-correlation matrix is computed with respect to the mapping epoch t_{k+1}:

$$\bar{\boldsymbol{P}}_{k+1}^{xy} = \sum_{j=0}^{2M} W_s^{(j)} (\bar{\boldsymbol{\chi}}_{j,k+1} - \hat{\boldsymbol{x}}_{k+1})(\tilde{\boldsymbol{\gamma}}_{j,k} - \bar{\boldsymbol{y}}_k)^T. \quad [20]$$

The mapping Kalman gain is computed as in Eq. [16]

$$\boldsymbol{K}_{k+1} = \bar{\boldsymbol{P}}_{k+1}^{xy} (\bar{\boldsymbol{P}}_k^{yy})^{-1} \quad [21]$$

where again an LU / least squares approach must be used. The *a posteriori* parameter covariance is then given by

$$\bar{\boldsymbol{P}}_{k+1}^{+} = \bar{\boldsymbol{P}}_{k+1}^{-} - \boldsymbol{K}_{k+1} \bar{\boldsymbol{P}}_{k+1}^{xy}. \quad [22]$$

This parameter / covariance mapping is not strictly required for the operation of the UKF and can therefore be performed as a post-processing step after the filter solution has been iterated to convergence. Note also that the UKF assumes Gaussian distributions for the sigma points, measurement estimates, and all covariances; while it accurately captures means and normal distribution values, higher order information is still lost.

III. Mars Approach

The first use case we examine is approach to Mars from deep space, either for orbit insertion or entry and landing. We first pose the simulation conditions then discuss results for two different scenarios defined by their initial epoch prior to Mars entry, $E-45$ and $E-5$ days. The only measurement type used is two-way Doppler tracking which is assumed to be supplied by the Deep Space Network (DSN).

III.i Simulation Set-Up

To assess the performance of the different filter implementations, we include a mix of bias parameters in addition to the dynamic state (position and velocity). In terms of spacecraft operations, we wish to estimate maneuver errors, the effect of solar radiation pressure (modeling the spacecraft as a sphere with nominal radius 1.5 m), uncertainty in the gravitational parameters of massive bodies (Sun, Mars, Earth), and errors in the location of ground tracking stations, as summarized in Table 1. Note that we present uncertainties in the gravitational parameters as fractions of their nominal value. The nominal states and maneuvers for the two different initial epochs are presented in Table 2, where all components are expressed in J2000 coordinates. We consider a variety of tracking schedules and initial errors in the position and velocity; while not all possible combinations are explored, the cases we consider include:

Table 1: Errors for Mars approach, all values 1-σ.

Parameter	**Error**	**Units**
Maneuver component	1.0	cm/s
SC radius	5	cm
Sun, Mars, & Earth μ	1e-9	–
Ground station position	1.0	cm
Two-way Doppler noise	5.62×10^{-3}	Hz

Table 2: Nominal states & maneuvers, Mars approach

Parameter	**Values**	**Units**
$E-5$ days		
Position	(2.01e8, 6.13e7, 2.26e7)	km
Velocity	(-5.41, 20.20, 9.60)	km/s
State epoch	21-NOV-2018 19:41:00 ET	–
Burn	(2, 1, -3)	cm/s
Burn epoch	24-NOV-2018 21:41:00 ET	–
$E-45$ days		
Position	(2.02e8, -1.17e7, -1.15e7)	km
Velocity	(5.13, 21.47, 9.87)	km/s
State epoch	12-OCT-2018 19:41:00 ET	–
Burn	(20, -10, 50)	cm/s
Burn epoch	10-NOV-2018 18:01:00 ET	–

- tracking pass lengths of 2-, 4-, and 6-hours;
- tracking passes every 1-, 3-, and 7-days;
- position errors ranging from 50- to 1000-km;

- velocity errors ranging from 50-cm/s to 10-m/s.

For all cases, the initial parameter covariances $\boldsymbol{P}$ are diagonal matrices whose non-zero elements are set equal to the squares of the 1-σ error values. When using the EKF, the tracking schedules are divided into 60-measurement sub-batches.

III.ii Results: E-45 days

We begin by analyzing the approach case where the spacecraft is 45 days away from Mars entry; measurements are taken over a span of 32 days with the goal of gaining an improved estimate of the spacecraft state prior to final approach. For this case we expect relatively benign dynamics up to the final few days, therefore the assumption of linearity will likely be valid along most of the reference path. We examine several different tracking strategies for a scenario where the spacecraft initial state is randomly perturbed in position and velocity with 1-σ values of 100 km and 1 m/s in each direction. The results for the LKF, EKF, and UKF are tabulated in Tables 3 and 4, where an empty line indicates failure to converge. We introduce errors in the initial state and velocity, the maneuver execution, as well as the parameters listed in Table 1; while we only report the converged RMS position error in the initial state, we do actively estimate all parameters.

Table 3: Results for approach E-45 days, varying pass length; initial pos. error 100-km, initial vel. error 1-m/s, tracking pass every 3 days.

Pass Dur.	**Pos. Err. (km, RMS)**	**Iter.**	**Time per iter. (s)**
Linear			
2 hrs	84.6	9*	0.7
4 hrs	71.2	9*	1.1
6 hrs	48.7	8*	1.4
Extended			
2 hrs	57.6	3*	0.9
4 hrs	39.9	4	1.5
6 hrs	15.8	5	2.1
Unscented			
2 hrs	55.3	5	8.6
4 hrs	14.4	3	19.7
6 hrs	9.6	3	39.4

*Case exhibits 'decaying walk'

In general, we see that the EKF and UKF outperform the LKF when the length of the tracking passes is varied (Table 3); the outcome is more mixed when considering the frequency of tracking (Table 4). Note that each case in Table 3 uses the same initial perturbation, while the cases in Table 4 also share a common initial condition that is different from the case in Table 3. The EKF and UKF are both able to converge when tracking occurs only once every 7 days, while the LKF does not. However, when it does converge, the LKF outperforms the UKF for this particular set of parameters perturbations. Note also that the computation times reported for the UKF are without parallelization.

Table 4: Results for approach E-45 days, varying intervals between passes; initial pos. error 100-km, initial vel. error 1-m/s, pass duration 4-hrs.

Pass Intv.	**Pos. Err. (km, RMS)**	**Iter.**	**Time per iter. (s)**
Linear			
7 days	–	–	–
3 days	21.4	4	1.0
1 day	18.3	1*	2.3
Extended			
7 days	36.1	9	0.9
3 days	26.6	6	1.5
1 day	6.7	2*	3.6
Unscented			
7 days	18.9	9	9.2
3 days	24.3	4	19.8
1 day	22.4	2	176.3

*Case exhibits 'decaying walk'

Some cases do exhibit a phenomenon we term a "decaying walk", that is, after reaching their lowest error value they begin steadily increasing in error while approaching some asymptotic limit. The magnitude of this limit seems to be correlated with the extent of tracking data available, perhaps related to a fundamental limit on the information content of those measurements.[25] To examine this behavior in more detail, we chart the per-iteration performance of a single case in Fig. 2, where this case case deliberately selected because it exhibits another interesting behavior sometimes exhibited by the batch UKF / EKF. This "jitter" occurs after convergence of the filter, where the RMS error increases then returns again to a lower level. Note that this behavior is not necessarily intrinsic to the UKF / EKF but could instead be caused by other factors in the simulation; more investigation is required to determine how this jitter could be mitigated, when it appears at all. Note that not all cases exhibit the phenomena illustrated

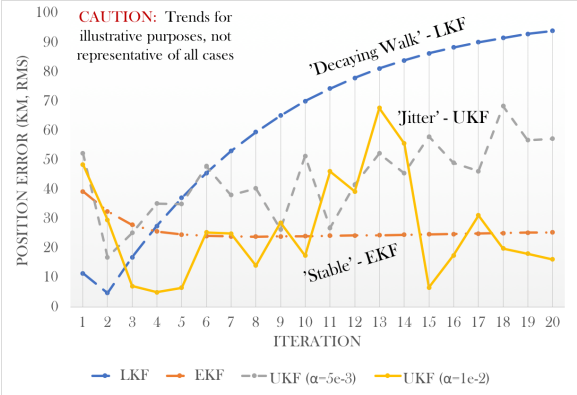

Fig. 2: Sample iteration histories for approach E-45 days; initial pos. error 100-km, initial vel. error 1-m/s, 4-hr tracking pass every 3 days.

in Fig. 2, but operators should be wary to spot them as they occur.

We now examine the performance of the batch filters subject to differing errors in the initial state, from large disturbances characteristic of worst-case upset recoveries down to error levels expected for routine operation. Results for the three filters are presented in Table 5; interestingly, the LKF and EKF exhibited instability for the smallest error case but the UKF was able to successfully converge. Note also that the LKF again exhibited a tendency to "walk" for this nominal approach case. At the moment, there seems to be no clear advantage to using the UKF or the EKF when operations are "routine", that is, relatively frequent tracking with sufficiently long passes. However, the two modified filters can be successfully used for cases with sparse tracking, a potentially useful feature for missions.

Before proceeding to the remaining analysis scenario, we first wish to examine the ability for each filter implementation to provide forward predictions in state uncertainty, e.g., covariance mapping. Position uncertainty ellipsoids for a sample case are depicted in Fig. 3, where the uncertainty is mapped forward to shortly before entry into the Mars system. As can be seen, the UKF provides a tighter uncertainty estimate than the LKF or EKF, while the EKF seems to perform worst of all. While is it hard to observe in the image, the uncertainty ellipsoids do also differ in orientation as well, though they are not wildly divergent in behavior. In theory, the UKF should provide the most accurate assessment of uncertainty, however a rigorous Monte Carlo analysis would be needed to confirm this assumption.

Table 5: Results for approach E-45 days, varying initial state errors; passes every 3 days, pass duration 4-hrs.

State Error	Pos. Err. (km, RMS)	Iter.	Time per iter. (s)
Linear			
1000 km 10 m/s	63.5	2*	0.8
500 km 5 m/s	14.8	3*	1.0
100 km 1 m/s	8.4	2*	1.1
50 km 0.5 m/s	–	–	–
Extended			
1000 km 10 m/s	56.3	2*	1.2
500 km 5 m/s	13.4	7	1.5
100 km 1 m/s	15.3	8	1.5
50 km 0.5 m/s	–	–	–
Unscented			
1000 km 10 m/s	22.0	10	14.2
500 km 5 m/s	1353	40**	20.3
100 km 1 m/s	12.8	7	19.9
50 km 0.5 m/s	30.2	8	19.3

*Case exhibits 'decaying walk'
**Case reached maximum iterations

III.iii Results: E-5 days

We turn now to the case of terminal approach to the Mars system, where our initial epoch is now 5 days prior to Martian entry. For this case, we consider only the variations in filter performance due to initial state errors. Recall that we only consider the use of two-way Doppler measurements for the sake of simplicity, although other measurement types such as ranging and DDOR would also likely be used to improve orbit determination accuracy. Regardless, the relative performance of the different filtering algorithms for this scenario are presented in Table 6. As can be seen, the UKF is the most accurate method by roughly a factor of 2, likely because it is able to more effectively handle the increased nonlinear dy-

namics due to the proximity to Mars. This opens the possibility that entry conditions can be targeted more precisely by switching the filtering algorithm, in addition to the acquisition of more precise measurements, as illustrated by the forward mappings shown in Fig. 4.

Table 6: Results for approach E-5 days, varying initial state errors; passes every 2 days, pass duration 6-hrs.

State Error	Pos. Err. (km, RMS)	Iter.	Time per iter. (s)
Linear			
500 km 5 m/s	59.4	3	0.5
100 km 1 m/s	56.7	4	0.5
50 km 0.5 m/s	60.5	4	0.5
Extended			
500 km 5 m/s	57.0	4	0.8
100 km 1 m/s	58.3	4	0.8
50 km 0.5 m/s	59.1	5	0.8
Unscented			
500 km 5 m/s	37.4	7	7.3
100 km 1 m/s	21.8	6	7.3
50 km 0.5 m/s	22.0	3	7.3

IV. Mars Orbit

We now consider the case when the spacecraft is in a Mars-centered orbit; as before, we first discuss the analysis set-up then the results for high- and low-altitude orbits. As with the approach cases, our only measurement type is DSN two-way Doppler tracking.

IV.i Simulation Set-Up

As with the approach case, we consider several bias parameters in addition to the dynamic position and velocity, as indicated in Table 7. We now omit the gravity of Earth, but include the Martian moons Phobos and Deimos; recall that the uncertainty in the gravitational parameters μ is set as a fraction of their true values. Two different orbits are con-

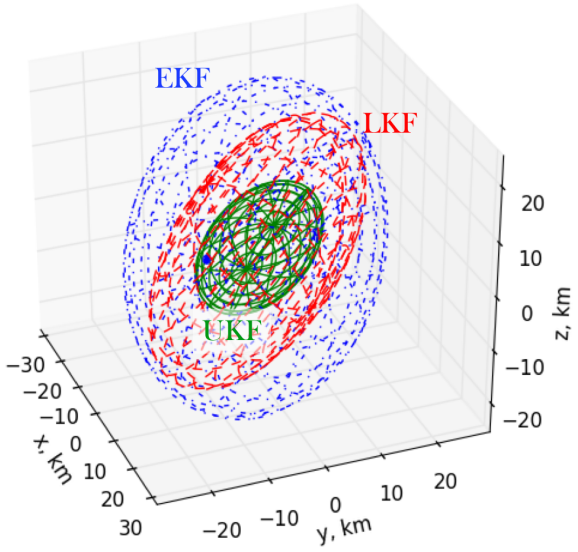

Fig. 3: Sample mapped covariance for approach E-45 days; initial pos. error 100-km, initial vel. error 1-m/s, 4-hr tracking pass every 3 days.

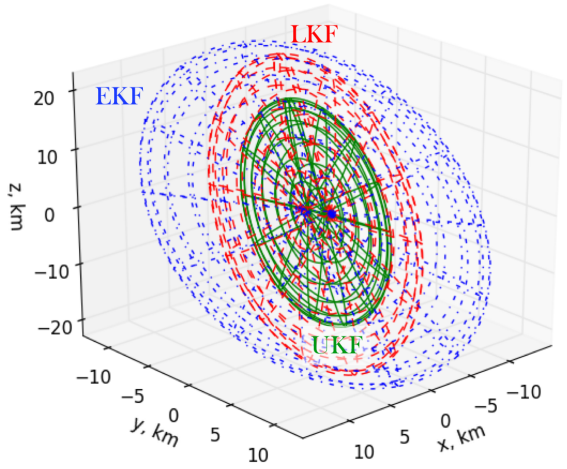

Fig. 4: Sample mapped covariance for approach E-5 days; initial pos. error 100-km, initial vel. error 1-m/s, 8-hr tracking pass every day.

Table 7: Errors for Mars orbit, all values 1-σ.

Parameter	Error	Units
Maneuver component	1.0	mm/s
SC radius	5	cm
Sun & Mars μ	1e-9	–
Phobos & Deimos μ	1e-4	–
Ground station position	1.0	cm
Two-way Doppler noise	5.62×10^{-3}	Hz

sidered, high- and low-altitude Mars orbits (HMO and LMO, respectively), where both are polar (inclination $\approx 95.4°$) and nearly circular (eccentricity ≈ 0.02). For both orbits, two maneuvers are executed at 6 hours and 12 hours after the initial epoch, with the burns being identical for both cases; tabulations of the initial states, maneuvers, and other orbital information are contained in Table 8. Position, velocity, and the maneuver components are all defined with respect to the Mars inertial equatorial reference frame. Note that the LMO is close enough to Mars that the atmosphere and its time variations would be a significant perturbing effect, however we have omitted stochastic drag from our investigation for simplicity and consistency between the HMO and LMO scenarios. As with the approach case, we consider a variety of position and velocity errors as well as tracking schemes, including:

Table 8: Nominal states & maneuvers, Mars orbit

Parameter	Values	Units
High Mars Orbit (HMO)		
Position	(5.14e3, 1.23e3, -5.35e3)	km
Velocity	(-1.58, -0.70, -1.64)	km/s
State epoch	01-Apr-2018 12:00:00 ET	–
Altitude	4072	km
Period	5.44	hrs
Low Mars Orbit (LMO)		
Position	(2.50e3, 0.60e3, -2.60e3)	km
Velocity	(-2.26, -1.00, -2.36)	km/s
State epoch	01-Apr-2018 12:00:00 ET	–
Altitude	233	km
Period	1.84	hrs
Maneuvers		
Burn #1	(-1, -2, 4)	cm/s
Epoch #1	01-APR-2018 18:00:00 ET	–
Burn #2	(11, -1, 0)	cm/s
Epoch #2	02-APR-2018 00:00:00 ET	–

- continuous tracking from 0.5-23.5 hours after the initial epoch;
- 3 tracking passes within the first 24 hours but bracketing the maneuvers, either 0.5 or 2 hours in duration apiece, as illustrated in Fig. 5;
- position errors ranging from 10-m to 50-km;
- velocity errors ranging from 1-cm/s to 5-m/s.

While these tracking schedules may not always be physically realizable due to occultations, we consider them as bounding cases to explore the performance of the different filtering schemes. For the EKF, the measurement set is divided into subsets of 30 measurements apiece.

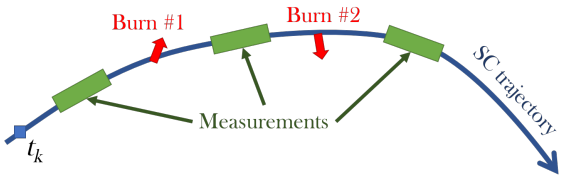

Fig. 5: Tracking scenario for the Mars-centric orbits.

IV.ii Results: High Mars Orbit (HMO)

We first consider the filter performances when applied to the case of orbit determination when the spacecraft is at a relatively high altitude. As can be seen in Table 9, the UKF performs comparably to or better than the other filters for sparser tracking. On the other hand, when tracking is continuous or very dense, the linear filters perform better. However, improved performance of the UKF is not always guaranteed, as we see when considering the results of various initial perturbations as highlighted in Table 10. While the UKF is able to converge for larger initial errors, there are cases where the EKF and even the LKF can offer more accurate estimates of the true state. This inconsistent pattern of performance could be partially related to the underlying

Table 9: Results for HMO, varying tracking schedule; initial pos. error 1-km, initial vel. error 10-cm/s.

Pass Dur.	Pos. Err. (m, RMS)	Iter.	Time per iter. (s)
Linear			
0.5 hrs	252	40**	0.9
2 hrs	232	4	1.0
Cont.	12.1	6	1.4
Extended			
0.5 hrs	167	40**	1.0
2 hrs	206	2*	1.3
Cont.	12.6	3	2.2
Unscented			
0.5 hrs	187	3	7.9
2 hrs	46.5	4	9.5
Cont.	21.5	9	16.0

*Case exhibits 'decaying walk'
**Case reached maximum iterations

spacecraft motion: while the orbit is Mars-centric, it is also high altitude with relatively long period so it is not dramatically nonlinear. Finally, we also examine the forward mapping of uncertainty, in this case two days forward from the initial epoch. Shown in Fig. 6, the uncertainty ellipses exhibit characteristic elongation in the along-track direction, though notably the UKF primary axis is roughly half that of the EKF/LKF. On the other hand, the intermediate and minimum axes are slightly larger for the UKF.

Table 10: Results for HMO, varying initial state errors; pass duration 2-hrs.

State Error	Pos. Err. (m, RMS)	Iter.	Time per iter. (s)
Linear			
50 km			
5 m/s	–	–	–
10 km			
1 m/s	–	–	–
1 km			
10 cm/s	0.6	21	1.0
100 m			
5 cm/s	95.0	2*	1.0
Extended			
50 km			
5 m/s	–	–	–
10 km			
1 m/s	272	40**	1.6
1 km			
10 cm/s	0.4	9	1.6
100 m			
5 cm/s	23.4	4*	1.6
Unscented			
50 km			
5 m/s	31.5	9	9.3
10 km			
1 m/s	632	10	9.6
1 km			
10 cm/s	66.2	6	9.4
100 m			
5 cm/s	144	3*	9.2

*Case exhibits 'decaying walk'

**Case reached maximum iterations

IV.iii Results: Low Mars Orbit (LMO)

Finally, we examine the scenario where the spacecraft is in low Mars orbit, with results summarized in Tables 11 and 12. In this case, the UKF performs comparably to or better than the other filters and again exhibits more robustness to initial errors. As noted before, these performance characteristics of the UKF could be exploited when tracking is expected to be sparse as well as during recovery from spacecraft upsets or safe modes. We also note that, in contrast to previous cases, the UKF now predicts larger uncertainties in forward predictions, as illustrated in Fig. 7. The majority of this increase is directed radially and in the cross-track direction, to the point where they are comparable or larger than the down-track axis; the uncertainty ellipsoids for the LKF and EKF, however, retain the traditional elongation in the velocity direction. Whether this increased uncertainty in the UKF results is a mathematical artifact or a true dynamic behavior will need to be confirmed via Monte Carlo or other statistical analyses.

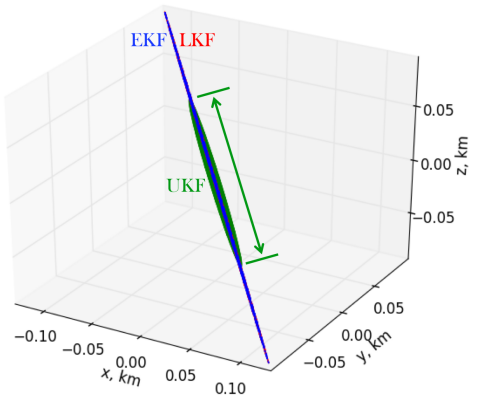

Fig. 6: Sample mapped covariance for HMO; initial pos. error 1-km, initial vel. error 10-cm/s, 2-hr tracking passes.

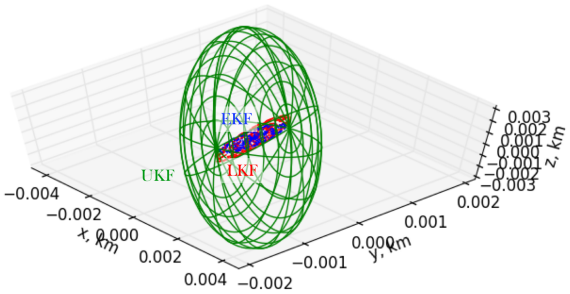

Fig. 7: Sample mapped covariance for LMO; initial pos. error 1-km, initial vel. error 10-cm/s, 2-hr tracking passes.

Table 11: Results for LMO, varying tracking schedule; initial pos. error 1-km, initial vel. error 10-cm/s.

Pass Dur.	Pos. Err. (m, RMS)	Iter.	Time per iter. (s)
Linear			
0.5 hrs	40.5	3*	7.8
2 hrs	1.0	4	8.3
Cont.	0.1	5	8.4
Extended			
0.5 hrs	128	3*	10.7
2 hrs	4.7	6	12.7
Cont.	0.1	4	19.6
Unscented			
0.5 hrs	24.8	3	40.1
2 hrs	1.9	6	40.8
Cont.	0.5	4	47.0

*Case exhibits 'decaying walk'

Table 12: Results for LMO, varying initial state errors; pass duration 2-hrs.

State Error	Pos. Err. (m, RMS)	Iter.	Time per iter. (s)
Linear			
10 km 1 m/s	–	–	–
1 km 10 cm/s	0.4	3	8.1
100 m 5 cm/s	1.2	4	8.7
Extended			
10 km 1 m/s	–	–	–
1 km 10 cm/s	3.5	7	12.6
100 m 5 cm/s	1.1	8	12.8
Unscented			
10 km 1 m/s	6.6	5	41.5
1 km 10 cm/s	1.9	3	42.3
100 m 5 cm/s	0.6	5	41.5

V. Conclusions

In this investigation, we applied the batched linear, extended, and unscented Kalman filters (LKF, EKF, and UKF, respectively) to analyze robotic missions to Mars, specifically the approach to Mars and operations while in orbit. In general, the unscented filter outperforms the EKF and LKF when the spacecraft motion is highly nonlinear, when large initial errors in state are present, or when measurements are sparse. Likewise, the EKF exhibits better convergence properties in these cases than simple linear filter. On the other hand, for more routine cases the LKF works as well as, and sometimes even better than, the UKF and EKF. Thus, the LKF will usually suffice for traditional missions, however the UKF and EKF could be highly useful during cases of upset recovery or when tracking is expected to be limited, for example on SmallSat missions. We observed some interesting trends in the operations of the filters, including "decaying walks" and "jitter" in the position estimates; these phenomena are not always exhibited, but operators should be aware to the possibility that they will occur. It is possible that some of this behavior is due to other factors in the simulation, but more investigation is required to determine the root cause. Finally, forward covariance predictions are usually more compact for the UKF than the LKF and EKF, the one exception being low Mars orbit; assuming that the unscented transform does in fact offer better predictions of covariances, this offers interesting avenues for improving Mars entry conditions and on-orbit operations without the need for new measurement types.

There are several interesting avenues for future work related to the filters in general and their application to Mars in particular. First, Monte Carlo analyses of the uncertainty distributions are needed to confirm the accuracy of the covariance ellipsoids. Second, various improvements can be made to the simulation scenarios can be made, including the modeling of stochastic parameters and the ability to consider biases rather than fully update them. Since the UKF uses tuning weights and the EKF can freely set the size of measurement sub-sets, a rigorous assessment of their impact on filter performance is called for, particularly if the UKF and EKF are to be used in operations. Likewise, all cases assumed initial state position and velocity errors with normal distributions and the same standard deviation in all directions; modeling error distributions as ellipsoids or other more realistic shapes would provide a higher fidelity assessment of filter performance. Furthermore, we have examined the filter performance assuming

ground processing, but assessing them in the context of on-board operations could help to determine the required computational needs for spacecraft autonomy. Finally, this analysis could be expanded to include other types of non-linear filters, for example the particle filter.

VI. Acknowledgements

This research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

References

[1] R. Kalman, "A New Approach to Linear Filtering and Prediction Problems", *Journal of Basic Engineering*, Vol. 82, 1960, pg. 35-45.

[2] P. Zarchan and H. Musoff, *Fundamentals of Kalman Filtering*, 4^{th} Ed., American Institute of Aeronautics and Astronautics, Reston, Virginia, 2015.

[3] M. Grewal, A. Andrews, and C. Bartone, *Global Navigation Satellite Systems, Inertial Navigation, and Integration*, 3^{rd} Ed., Wiley, Hoboken, New Jersey, 2013.

[4] B. Tapley, B. Schutz, and G. Born, *Statistical Orbit Determination*, Elsevier, Inc., Burlington, Massachusetts, 2004.

[5] A. Jazwinski, *Stochasitc Processes and Filtering Theory*, Academic Press, New York, New York, 1970.

[6] B. Ristic, S. Arulampalam, and N. Gordon, *Beyond the Kalman Filter: Particle Filters for Tracking Applications*, Artec House, Boston, Massachusetts, 2004.

[7] S. Julier, J. Uhlmann, H. Durrant-Whyte, "A New Approach for Filtering and Nonlinear Systems", *Proceedings of the American Control Conference*, IEEE, New York, 1995, pg. 1628-1632.

[8] S. Julier, J. Uhlmann, H. Durrant-Whyte, "A New Method for the Nonlinear Transformation of Means and Covariances in Filters and Estimators", *IEEE Transactions on Automatic Control*, Vol. 45., No. 3, March 2000, pg. 477-482.

[9] S. Julier and J. Uhlmann, "Unscented Filtering and Nonlinear Estimation", *Proceedings of the IEEE*, Vol. 92, No. 3, March 2004, pg. 401-422.

[10] D-J. Lee and K. Alfriend, "Sigma Point Filtering for Sequential Orbit Estimation and Prediction", *Journal of Spacecraft and Rockets*, Vol. 44, No. 2, March-April 2007, pp. 388-398.

[11] J. Carpenter, S. Hur-Diaz, and F. Markley, "Generalized Covariance Analysis of Additive Divided-Difference Sigma-Point Filters", *Advances in the Astronautical Sciences*, Vol. 135, 2009, pg. 69-84.

[12] P. Pardal, H. Kuga, and R. Moraes, "Unscented Kalman Filter Robustness Assessment for Orbit Determination Using GPS Signals", *Advances in the Astronautical Sciences*, Vol. 148, 2013, pg. 117-134.

[13] X. Zhao, S. Liu, and C. Han, "Application of UKF in Autonomous Orbit Determination of Navigation Constellation", *Advances in the Astronautical Sciences*, Vol. 140, 2011, pg. 1757-1770.

[14] D. Lee and H. Pernicka, "Vision-Based Relative State Estimation using the Unscented Kalman Filter", *Advances in the Astronautical Sciences*, Vol. 134, 2009, pg. 1005-1026.

[15] C. Olson, R. Russell, and J. Carpenter, "Small Body Optical Navigation using the Additive Divided Difference Sigma Point Filter", *Advances in the Astronautical Sciences*, Vol. 152, 2014, pg. 3027-3045.

[16] R. Nakamura, C. Hirose, H. Ikeda, and K. Nakajima, "Real-Time Orbit Determination for Lunar and Planetary Missions", *Advances in the Astronautical Sciences*, Vol. 148, 2013, pg. 135-146.

[17] T. Ely and J. Seubert, "Batch Sequential Estimation with Non-Uniform Measurements and Non-Stationary Noise", *Advances in the Astronautical Sciences*, Vol. 160, 2017.

[18] ES. Park, SY. Park, KM. Roh, and KH. Choi, "Satellite Orbit Determination using a Batch Filter based on the Unscented Transformation", *Aerospace Science and Technology*, Vol. 14, 2010, pg. 387-396.

[19] S. Evans, et al. "MONTE: The Next Generation of Mission Design & Navigation Software", *The* 6^{th} *International Conference on Astrodynamics Tools and Techniques (ICATT)*, Darmstadt, Germany, March 14-17, 2016, `https://indico.esa.int/indico/event/111/session/30/contribution/177`

[20] Solar System Dynamics Group, "HORIZONS System", Jet Propulsion Laboratory, California Institute of Technology, `http://ssd.jpl.nasa.gov/?horizons`

[21] R. Bate, D. Mueller, and J. White, *Fundamentals of Astrodynamics*, Dover, Inc., New York, New York, 1971.

[22] D. Vallado and W. McClain, *Fundamentals of Astrodynamics and Applications*, 3^{rd} Ed., Microcosm / Springer, El Segundo, California, 2007.

[23] T. Moyer, *Formulation for Observed and Computed Values of Deep Space Network Data Types for Navigation*, John-Wiley & Sons, Inc., Hoboken, New Jersey, 2003.

[24] G. Bierman, *Factorization Methods for Discrete Sequential Estimation*, Dover, Inc., New York, New York, 1977.

[25] T. Hamilton and W. Melbourne, "Information Content of a Single Pass of Doppler Data from a Distant Spacecraft", Jet Propulsion Laboratory, California Institute of Technology, May 31, 1966, pg. 18-23, Report No. JPL-SPS 37-39, `http://descanso.jpl.nasa.gov/history/DSNTechRefs.html`